

US006889304B2

(12) **United States Patent**  
**Perego et al.**

(10) **Patent No.:** **US 6,889,304 B2**  
(45) **Date of Patent:** **May 3, 2005**

(54) **MEMORY DEVICE SUPPORTING A DYNAMICALLY CONFIGURABLE CORE ORGANIZATION**

(75) Inventors: **Richard E. Perego**, San Jose, CA (US); **Donald C. Stark**, Los Altos Hills, CA (US); **Frederick A. Ware**, Los Altos Hills, CA (US)

(73) Assignee: **Rambus Inc.**, Los Altos, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

5,787,267 A	*	7/1998	Leung et al. ....	711/105
5,793,998 A		8/1998	Copeland et al. ....	395/309
5,801,985 A		9/1998	Roohparvar et al. ...	365/185.01
5,852,725 A		12/1998	Yen .....	395/282
5,893,927 A		4/1999	Hovis .....	711/171
5,958,033 A		9/1999	Schubert et al. ....	710/126
6,034,878 A		3/2000	Osaka et al. ....	365/52
6,047,347 A		4/2000	Hansen et al. ....	710/127
6,125,157 A		9/2000	Donnelly et al. ....	375/371
6,138,185 A		10/2000	Nelson et al. ....	710/33
6,141,273 A		10/2000	Ku et al. ....	365/201
6,144,220 A		11/2000	Young .....	326/41
6,311,313 B1		10/2001	Camporese et al. ....	716/6
6,625,687 B1		9/2003	Halbert et al. ....	711/105

(21) Appl. No.: **10/302,420**

(22) Filed: **Nov. 22, 2002**

(65) **Prior Publication Data**

US 2004/0019756 A1 Jan. 29, 2004

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 09/797,099, filed on Feb. 28, 2001.

(60) Provisional application No. 60/392,098, filed on Jun. 26, 2002.

(51) **Int. Cl.**<sup>7</sup> ..... **G06F 12/00**; G06F 12/02

(52) **U.S. Cl.** ..... **711/170**; 711/5; 711/115

(58) **Field of Search** ..... 711/115, 170, 711/5

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,766,538 A	8/1988	Miyoshi .....	364/200
4,985,867 A	1/1991	Ishii et al. ....	365/221
5,394,528 A	2/1995	Kobayashi et al. ....	395/325
5,530,814 A	6/1996	Wong et al. ....	395/312
5,559,970 A	9/1996	Sharma .....	395/312
5,614,855 A	3/1997	Lee et al. ....	327/158
5,652,870 A	7/1997	Yamasaki et al. ....	395/500
5,717,871 A	2/1998	Hsieh et al. ....	395/250
5,717,901 A	2/1998	Sung et al. ....	395/497.01

**FOREIGN PATENT DOCUMENTS**

EP	0887737 A2	10/1997	.....	G06F/13/40
GB	2367400 A	4/2001	.....	G06F/12/06

**OTHER PUBLICATIONS**

The Authoritative Dictionary of IEEE Standards Terms, Copyright 2000, IEEE Press, Seventh Edition, p. 252.\*

(Continued)

*Primary Examiner*—Mano Padmanabhan

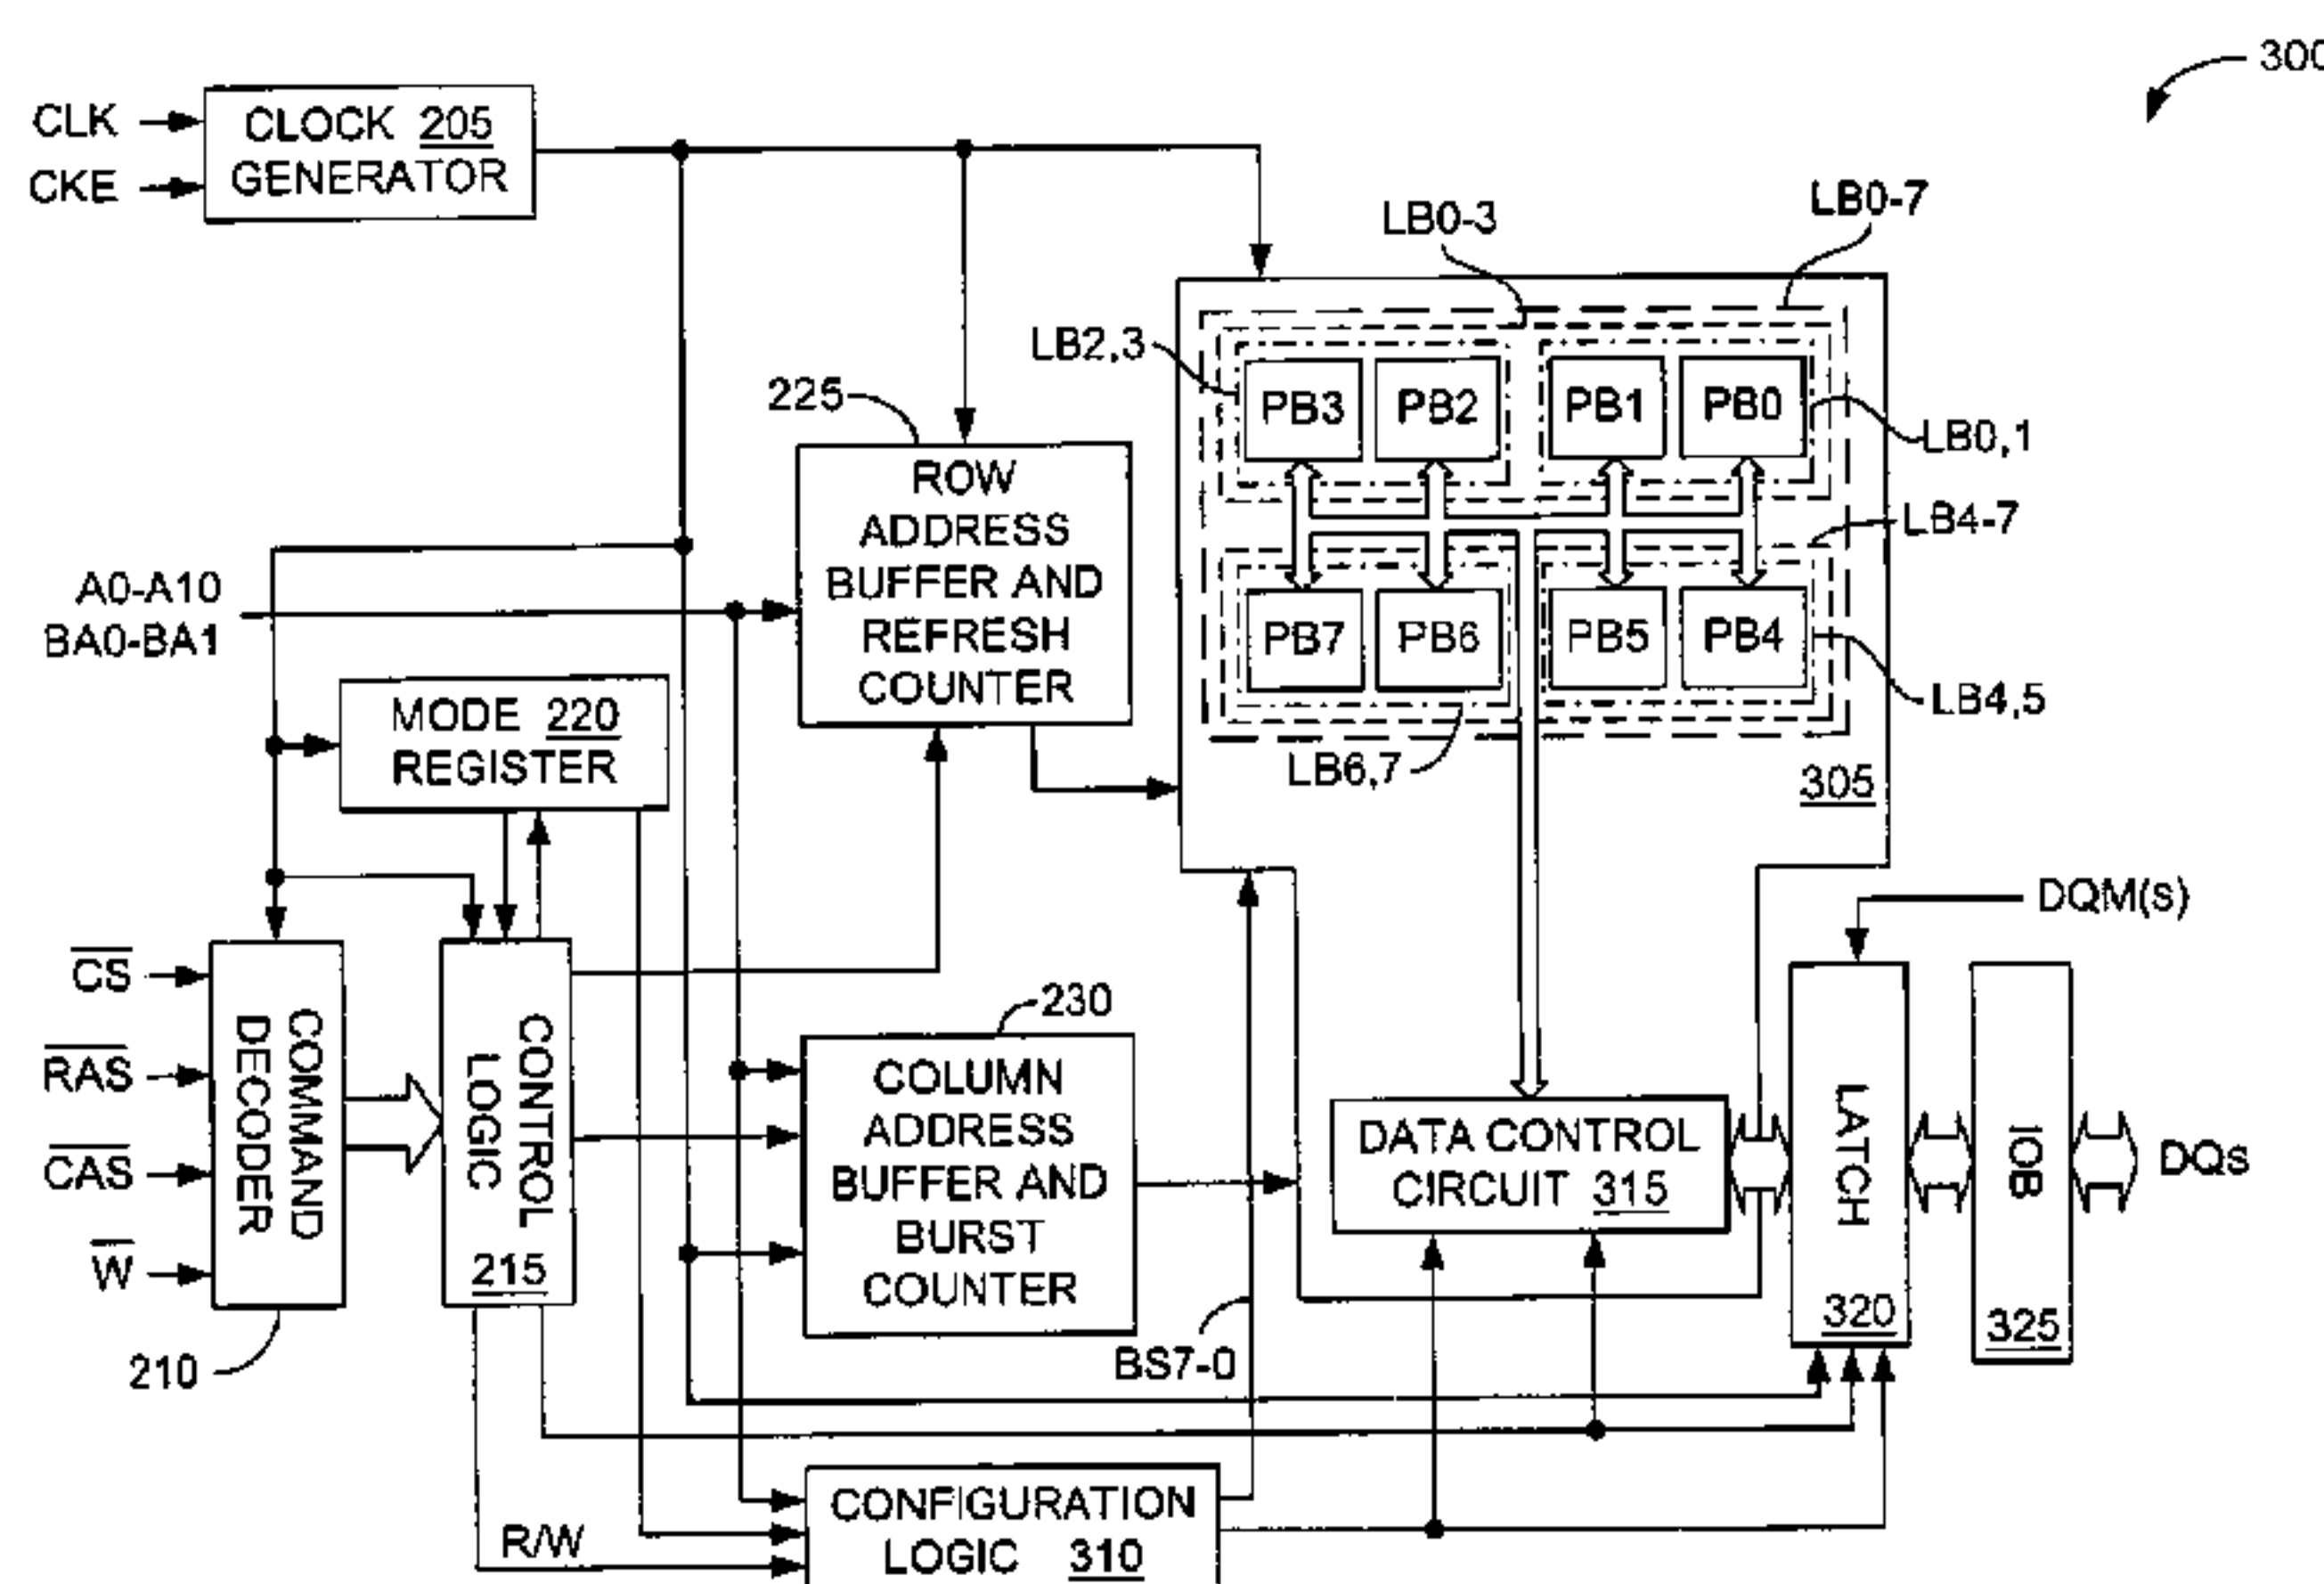
*Assistant Examiner*—Midys Inoa

(74) *Attorney, Agent, or Firm*—Silicon Edge Law Group LLP; Arthur J. Behiel

(57) **ABSTRACT**

Described is a memory system in which the memory core organization changes with device width. The number of physical memory banks accessed reduces with device width, resulting in reduced power usage for relatively narrow memory configurations. Increasing the number of logic memory banks for narrow memory widths reduces the likelihood of bank conflicts, and consequently improves speed performance.

**14 Claims, 15 Drawing Sheets**



OTHER PUBLICATIONS

“A 390–mm<sup>2</sup>, 16–Bank, 1–Gb DDR SDRAM with Hybrid Bitline Architecture”, Kirihata, et al. IEEE Journal of Solid–State Circuits, vol. 34, No. 11, Nov. 1999. pp. 1580–1588.


TA 11.6 “An 18Mb, 12.3GB/s CMOS Pipeline–Burst Cache SRAM with 1.54Gb/pin”, Zhao, et al. IEEE International Solid–State Circuits Conference, 1999. 10 pages.

“A 209K–Transistor ECL Gate Array with RAM”, Satoh, et al. IEEE Journal of Solid–State Circuits, vol., 24. No. 5, Oct. 1989. pp. 1275–1281.

“Configurable On–Chip RAM Incorporated Into High Speed Logic Array”, Rodney T. Masumoto. Proceedings of the IEEE 1985 Custom Integrated Circuits Conference, May 20–23, 1985. pp. 243–245.

“IEEE 100 The Authoritative Dictionary of IEEE Standards Terms.” Seventh Edition. Critical piece first to Cross bar switch (p. 252).

\* cited by examiner



**FIG. 1**  
(PRIOR ART)

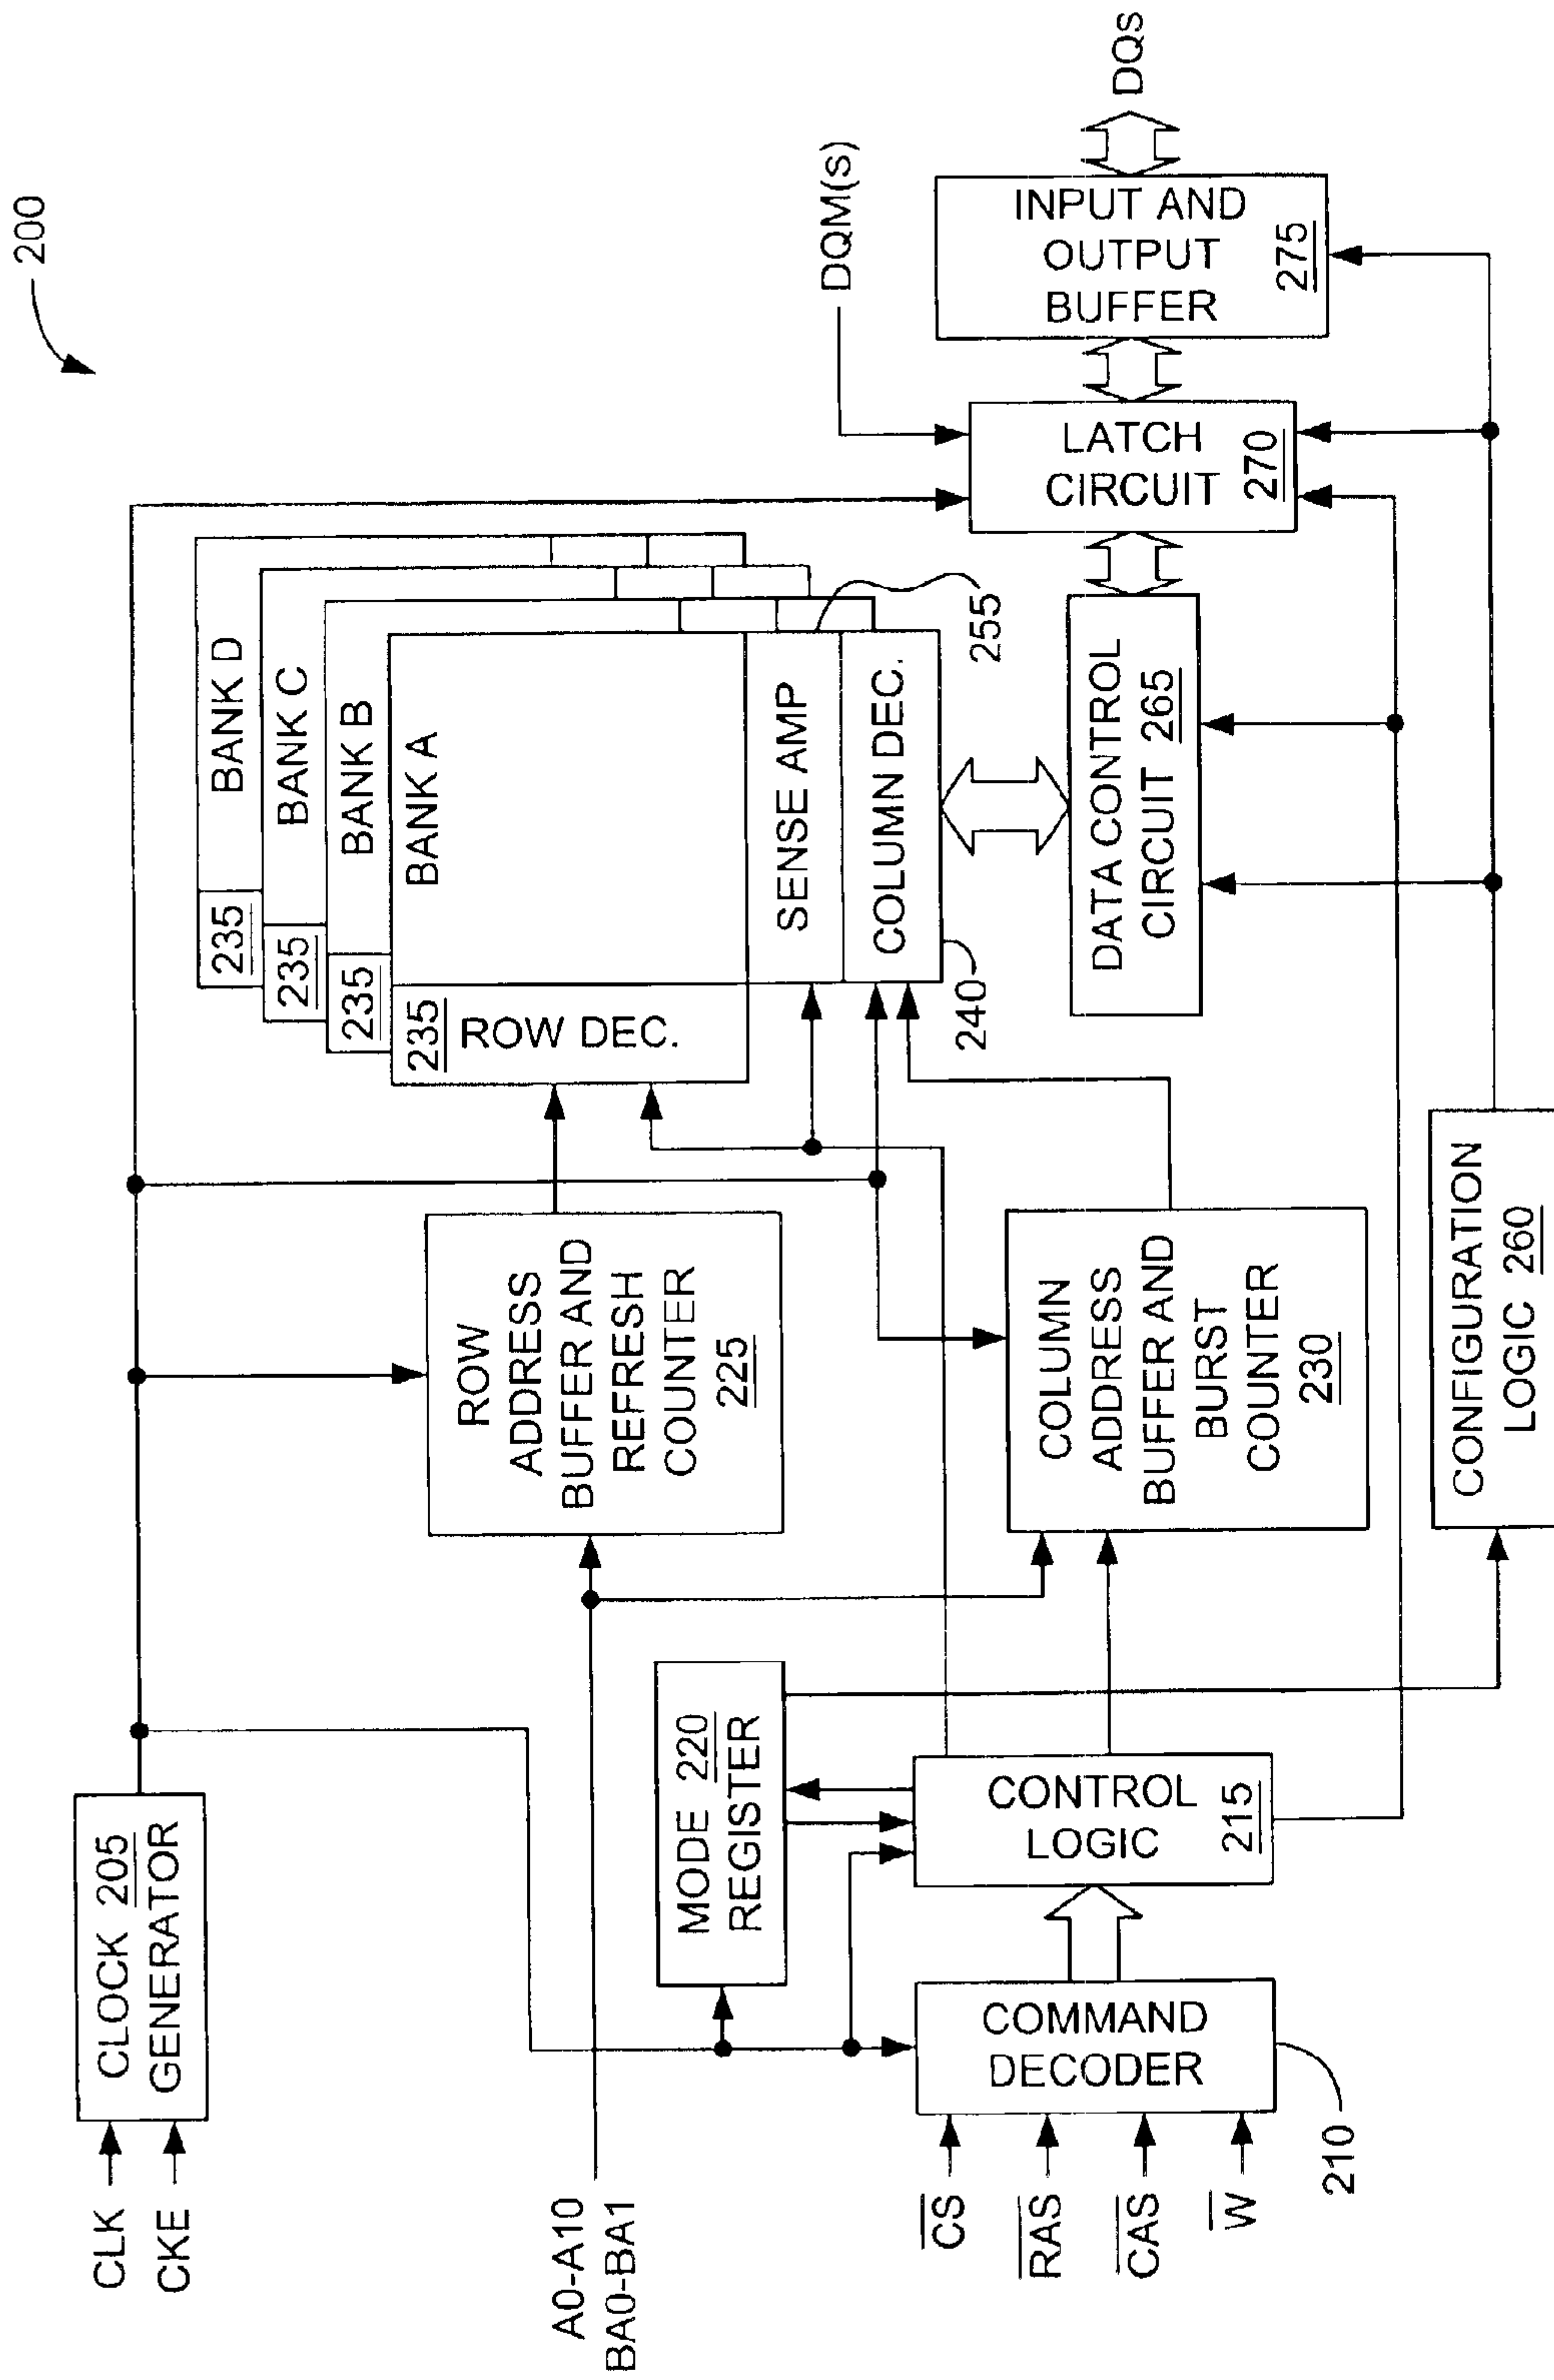


FIG. 2 (PRIOR ART)



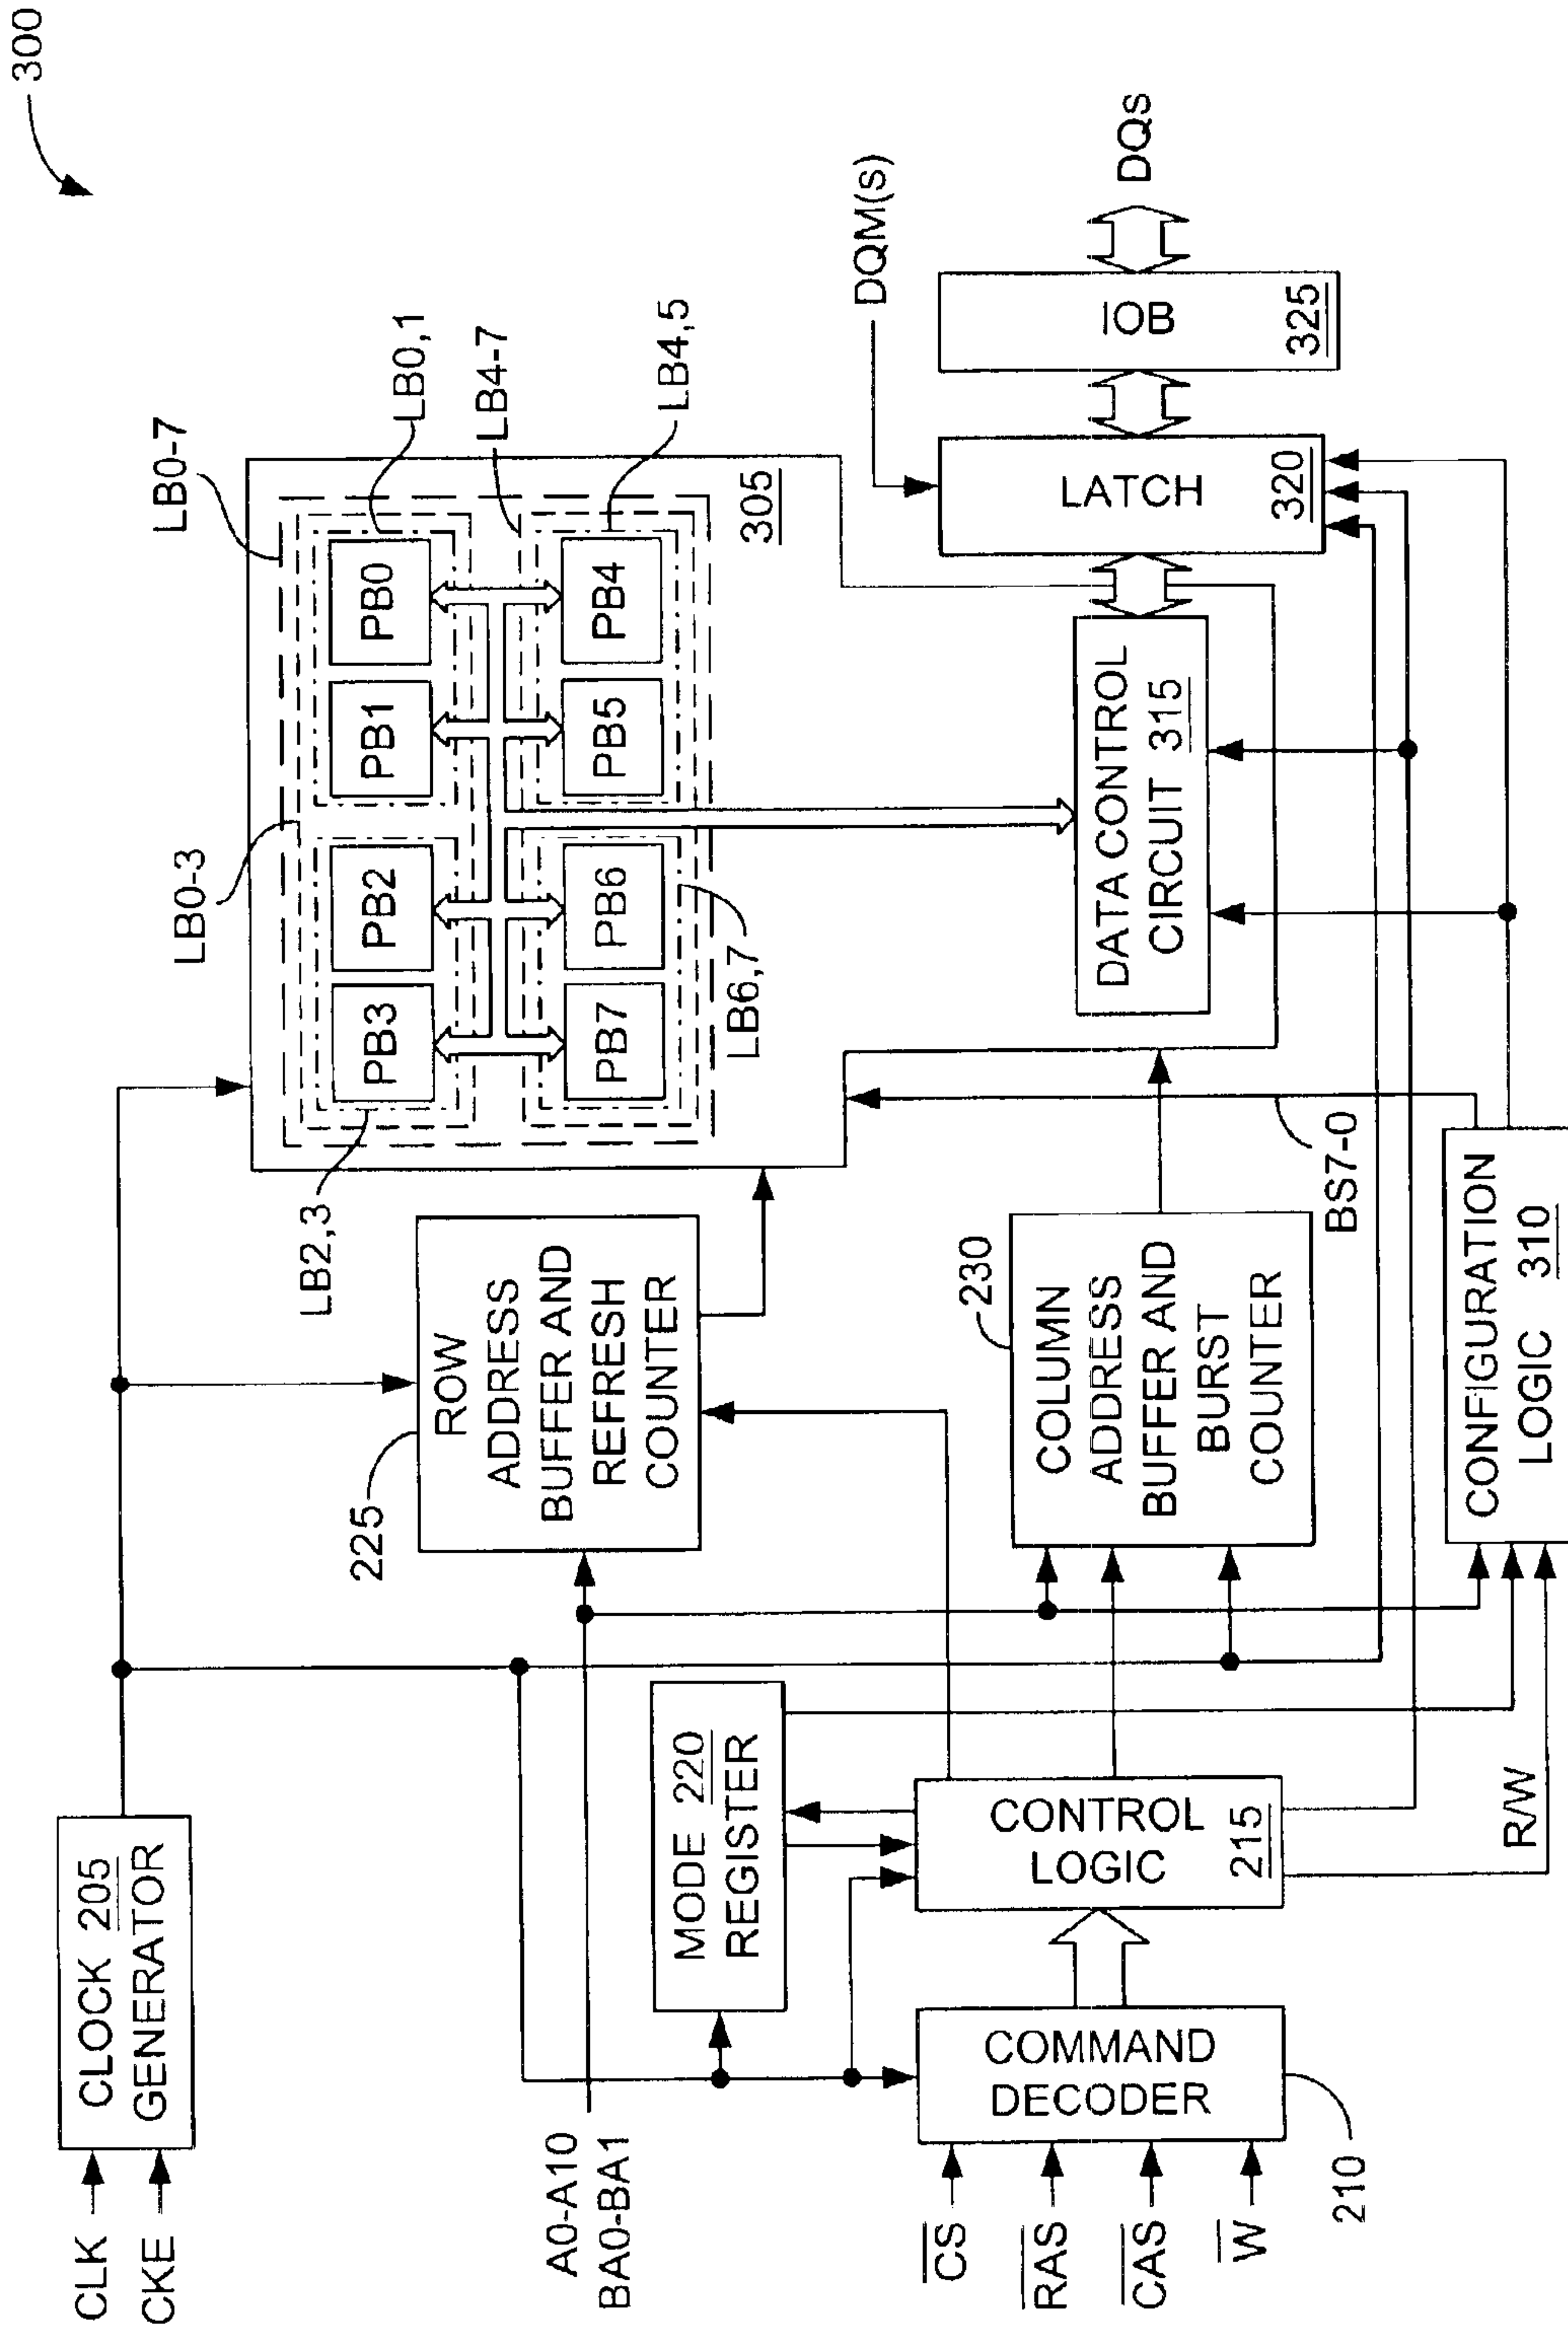


FIG. 3

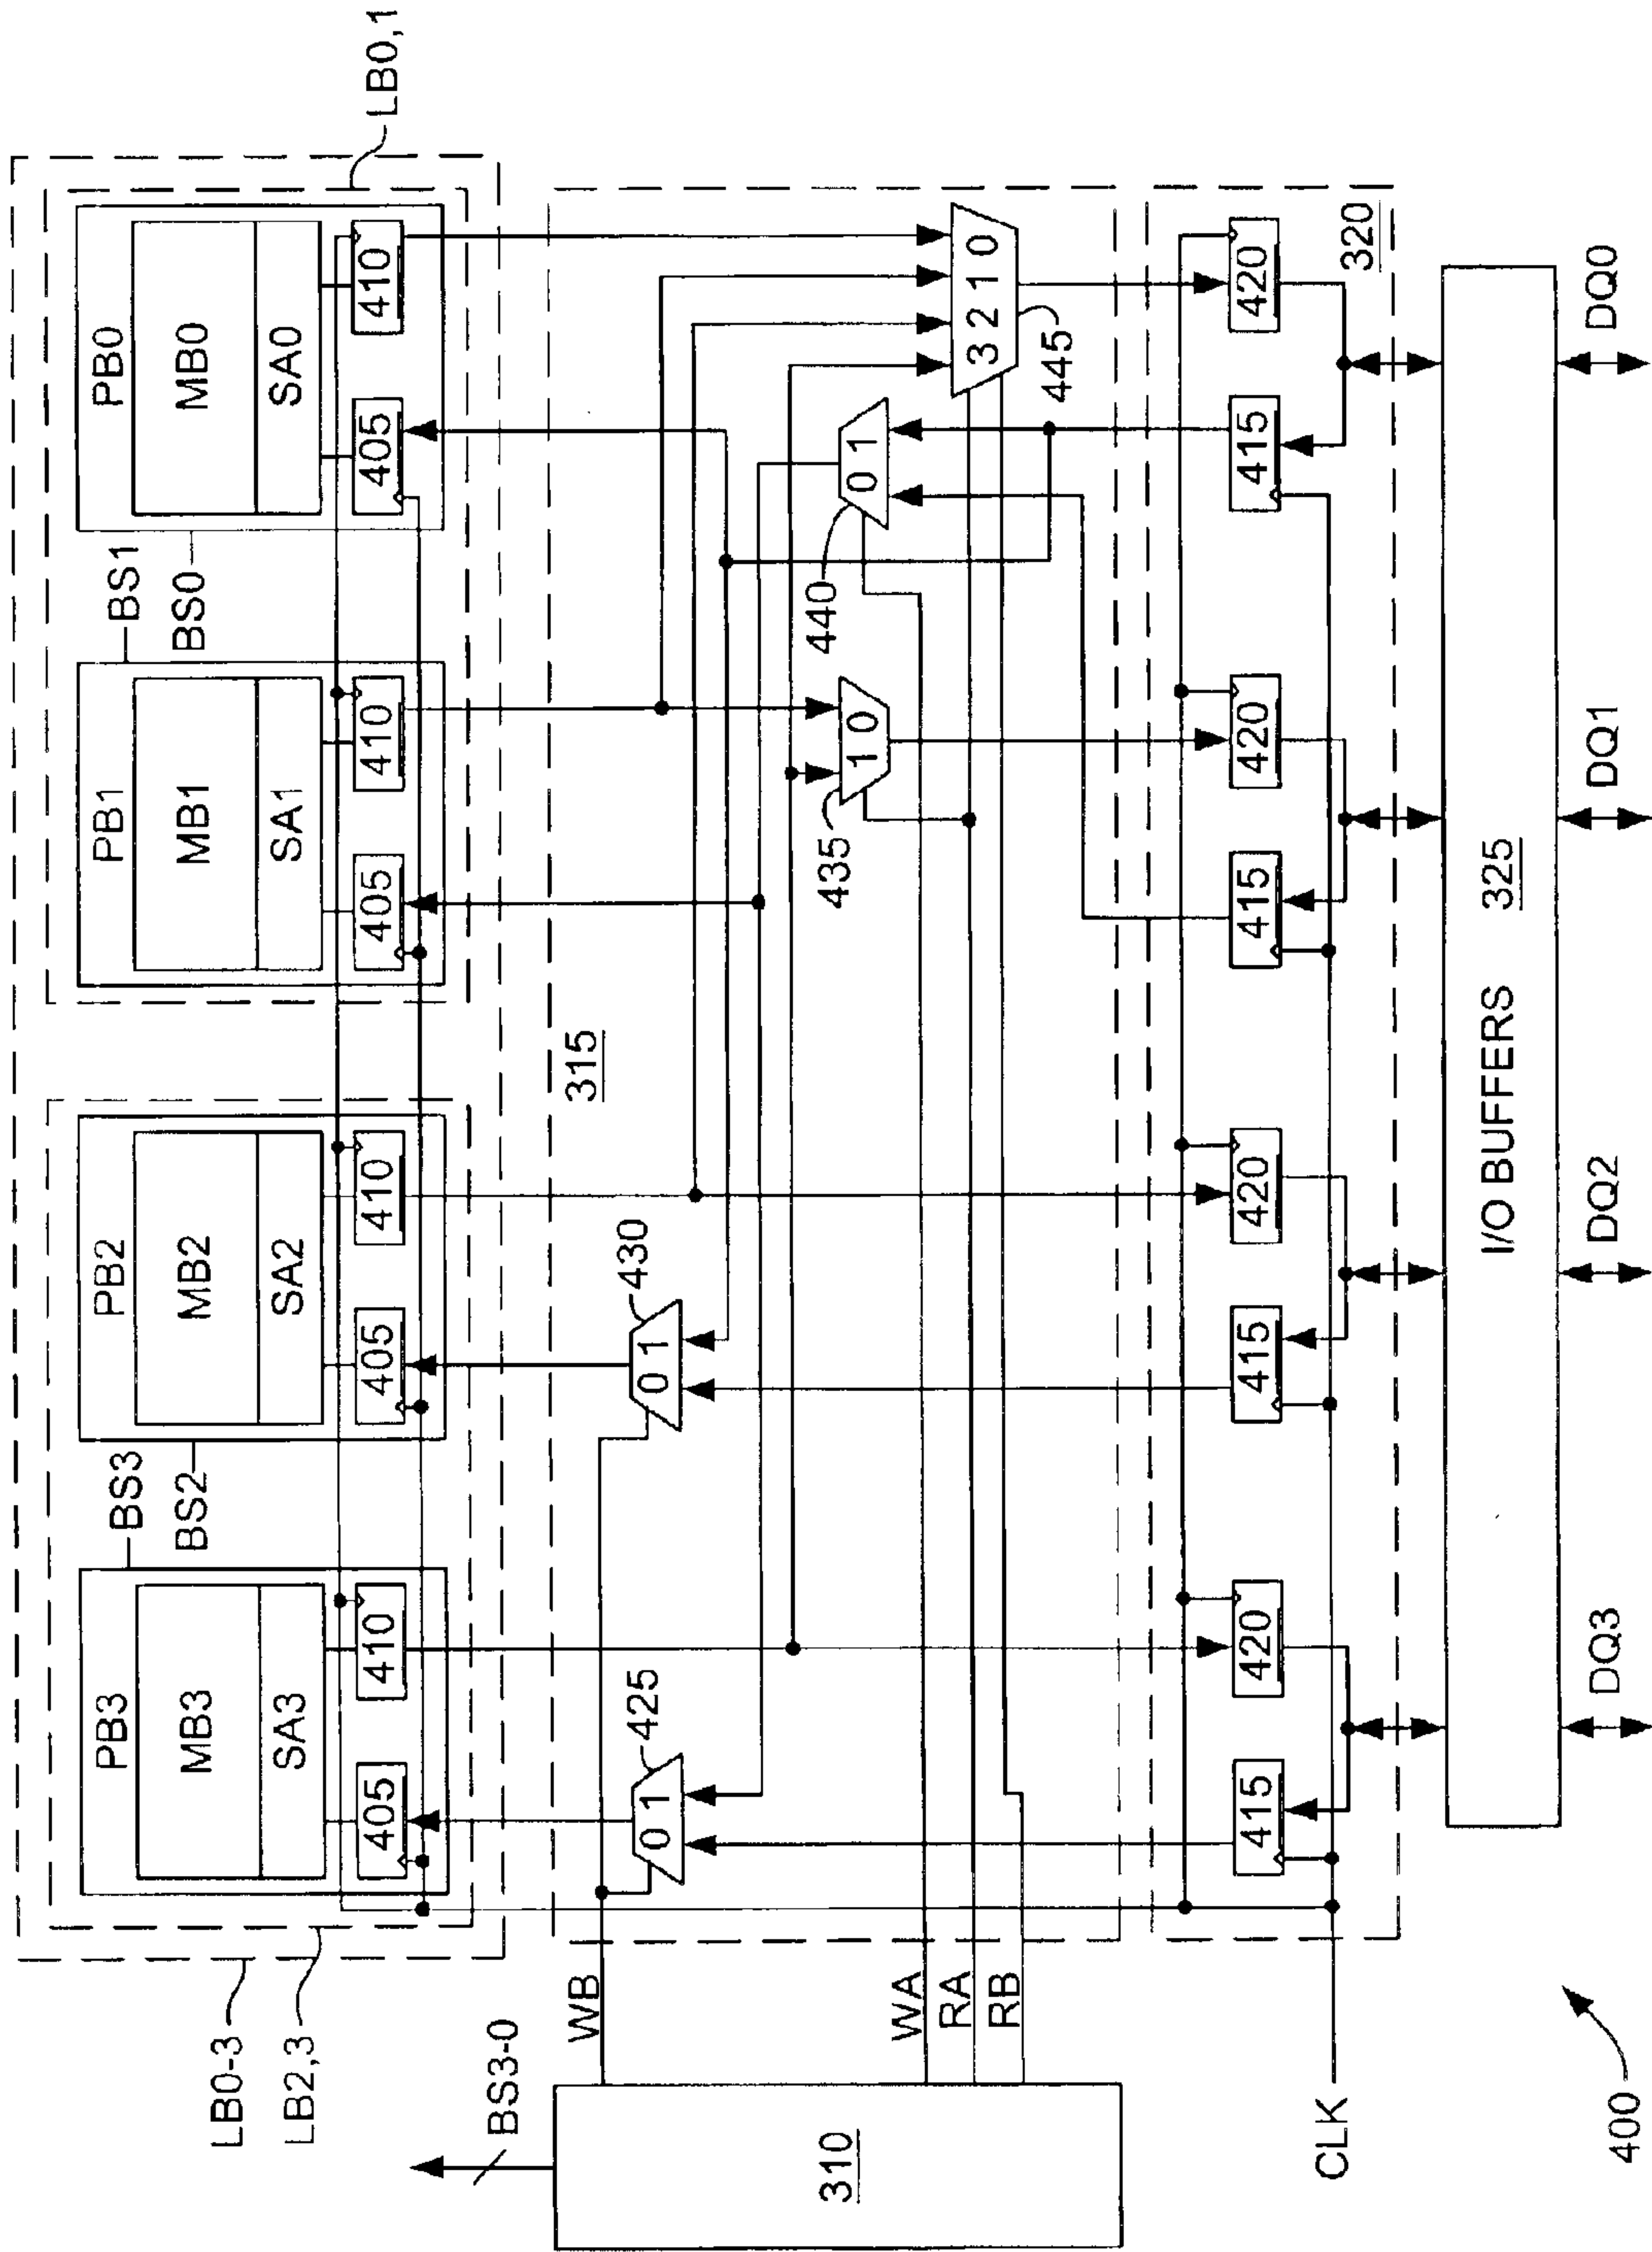


FIG. 4A




FIG. 4B




FIG. 5A




FIG. 5B




FIG. 5C






FIG. 6A




FIG. 6B




FIG. 7




FIG. 8




FIG. 9A




FIG. 9B




FIG. 9C




FIG. 9D




FIG. 10 (PRIOR ART)






FIG. 11A




FIG. 11B






FIG. 12




FIG. 13A




FIG. 13B




FIG. 13C




FIG. 13D




FIG. 13E




FIG. 13F



## MEMORY DEVICE SUPPORTING A DYNAMICALLY CONFIGURABLE CORE ORGANIZATION

### BACKGROUND

FIG. 1 (prior art) shows an example of a conventional memory system **100**. In this example, memory system **100** resides on a computer motherboard **105** and is actually a subsystem of the motherboard. System **100** includes a plurality of female electrical connectors **110**, each of which accepts a memory module **115** (only one of which is shown here). Each memory module **115** contains a plurality of memory devices **120**, typically packaged as discrete integrated circuits (ICs). Memory devices **120** are usually some type of read/write memory, such as Dynamic Random Access Memories (DRAMs), Static Random Access Memories (SDRAMs), Flash RAM, or other types. Read-Only Memories (ROM) devices might also be used.

Motherboard **105** includes a memory controller **125** connected via conductive traces **130** to connectors **110**. Memory controller **125** communicates with memory modules **115** through conductive traces **130**. Memory controller **125** also has an interface (not shown) that communicates with other components on the motherboard, allowing those components to read from and write to memory.

Each memory module **115** typically contains a fixed-width data path interface. The fixed-width nature of the interface is generally a result of a desire to create an industry standard interface that can accommodate interoperable modules from a large number of suppliers. System **100** works with different numbers of memory modules **115** installed, and with modules having different memory capacities and/or organizations. However, a system such as this is normally designed for a specific system data path width, i.e., for a specified number of data bit lines from controller **125** to memory modules **115**.

Memory devices can be targeted to a wide variety of markets with very different sets of cost and performance constraints; consequently, the optimal device width can vary significantly from one application to the next. Unfortunately, these variations make it difficult for memory suppliers and distributors to accurately predict the customer demand mix for memory devices of various widths. Inaccuracies in demand-mix prediction can cause supply/demand imbalances and inventory management difficulties, which in turn can lead to pricing instability and highly variable profit margins. Furthermore, a memory device manufacturer may find that optimizing the cost for each target device width means a different design at the die level and potentially at the package level. This can increase the time-to-market and level of financial and engineering resources required to deliver each of these products to market.

Fixed-width devices have other drawbacks related to inflexible data path configuration. Because the system memory interface width and memory device interface widths are fixed, the addition of more memory devices or modules to the system typically requires multiple ranks, which generally necessitates the use of a multi-drop datapath topology. Adding more drops to the system tends to degrade signaling performance.

One way to reduce time-to-market and resource requirements is to create a common die design and package pinout that can support a variety of device data path widths. Some memory manufacturers support this capability through memory designs that allow configurations to be postponed

until relatively late in the manufacturing process. A configuration is typically selected through one of several possible schemes, such as fuse or anti-fuse programmability, wire-bonding options, or upper level metal mask changes. This flexibility allows the device to be tested at the target width and sold as a fixed-width device.

Another way to reduce time-to-market and resource requirements associated with fixed-width memories is to use a memory design in which the width (e.g., the number of data pins) can be dynamically changed to suit the needs of a particular system. One such memory design is depicted in U.S. patent application No. 5,893,927 to William P. Hovis, which is incorporated herein by reference. FIG. 2, taken from the Hovis patent, illustrates a conventional synchronous dynamic random access memory (SDRAM) **200** having a programmable device width. SDRAM **200** includes a clock generator **205** that provides clock signals to various components of SDRAM **200**. A command decoder **210** receives chip select ICS, row enable /RAS, column enable /CAS and write command /W inputs. (The "/" preceding the signal names identifies the signals as active low. Overbars are used in the figures for the same purpose.) Command decoder **210** recognizes, for example, a write command when /CS, /CAS, and /W are simultaneously asserted (i.e., logic low). Command decoder **210** then outputs the command to some control logic **215**, which controls the operation of the other components of SDRAM **200** based on the received command.

Besides the commands of /CS, /RAS, /CAS and /W, command decoder **210** also recognizes commands based on a combination of /CS, /RAS, /CAS, and /W. For instance, command decoder **210** decodes the simultaneous receipt of /CS, /RAS, /CAS, and /W as a mode register set command. When a mode register set command is received, control logic **215** causes a mode register **220** to latch the address data on address inputs **A0–A10** and **BA0–BA1**.

The data on address inputs **A0–A10**, generally, represent either a row or column address, whereas the data on address inputs **BA0–BA1**, generally, represent a bank address. The bank addresses inputs **BA0–BA1** specify one of the memory banks **A–D** discussed in detail below. During the mode register set operation, however, the data on address inputs **A0–A10** and **BA0–BA1** represent commands. Hereinafter, the address inputs and the data thereon will generically be referred to as address inputs.

SDRAM **200** includes a row address buffer and refresh counter **225** and a column address buffer and burst counter **230**, both of which connect to address inputs **A0–A10** and **BA0–BA1**. The row address buffer portion latches the address inputs at row-access-strobe (RAS) time and provides the row address to the appropriate row decoder **235**. The refresh counter portion refreshes the memory. The column address buffer portion latches the address inputs at column-access-strobe (CAS) time and provides the column address to the appropriate column decoder **240**. The burst counter portion controls the reading/writing of more than one column based on a pre-set burst length.

The memory of SDRAM **200** is divided into four memory banks **A–D** that can be independently and simultaneously selected. Each memory bank **A–D** has associated therewith a row decoder **235**, a sense amplifier **255**, and a column decoder **240**. Based on the address latched by the row address buffer and refresh counter **225**, one of row decoders **235** enables a row of bits in the corresponding bank. An associated sense amplifier **255** latches the columns of this row via sense amplification, and the associated column



decoder **240** outputs one or more bits depending on the device width and burst length. Sense amplifier **255** typically represents a combination of column I/O amplifiers arranged along an edge of the array of banks and lower-level sense amplifiers interleaved between memory cells.

SDRAM **200** includes configuration logic **260** for setting the device width. Configuration logic **260** connects to mode register **220**, and from there receives a memory-width configuration value stored in register **220** during device configuration. Based on this information, configuration logic **260** configures a data control circuit **265**, a latch circuit **270**, and an input/output (I/O) buffer **275** to obtain the device width associated with the memory-width configuration value. Specifically, configuration logic **260** controls switches and multiplexers in data control circuit **265** such that the number of active I/O drivers corresponds to the programmed device width.

Data control circuit **265** is connected to each column decoder **240**, and to data I/O pin(s) DQ(s) via latch circuit **270** and input/output buffer **275**. During a read operation, sense amplifiers **255** and column decoders **240** output data to data control circuit **265** based on the row enabled by decoder **235**, the column enabled by decoder **240**, and the burst length. Data control circuit **265** then routes the data to the number of I/O drivers set based on the device width. The data from the I/O drivers is then latched by the latch circuit **270**, buffered by I/O buffer **275**, and output on the data I/O pin(s) DQ(s). The number of I/O pin(s) DQ(s) corresponds to the device width.

During a write operation, SDRAM **200** receives data over the I/O pin(s) DQ(s). This data is buffered by I/O buffer **275**, latched by latch circuit **270**, and received by data control circuit **265**. Data control circuit **265** sends the data to the appropriate column decoders **240** for storage in the memory banks A–D according to the enabled row and column.

SDRAM **200** also includes an input DQM to latch circuit **270** for every 8 bits of input/output. For instance a  $\times 16$  SDRAM will have two inputs DQM0 and DQM1. When enabled, the input DQM prevents reading or writing the remainder of a burst. In this manner, the burst length can be controlled.

Each read operation presents an entire row of data to sense amps **255**. Each write operation similarly involves an entire row. In SDRAM **200**, changing the memory width merely changes the number of bits selected from the accessed row: the narrower the memory configuration, the fewer bits are selected from the accessed row. Since the power required to perform a row access does not change with changes in device width, the relative power efficiency of row accesses reduces with memory width.

#### BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 (prior art) shows an example of a conventional memory system **100**.

FIG. 2 (prior art) illustrates a conventional synchronous dynamic random access memory (SDRAM) **200** having a programmable device width.

FIG. 3 depicts a variable-width memory **300** in accordance with an embodiment of the invention.

FIG. 4A details a portion of an embodiment of memory **300** of FIG. 3.

FIG. 4B details a portion of another embodiment memory **300** of FIG. 3.

FIGS. 5A–5C depict various width configurations of a memory module **500** that includes four variable-width

memories **300** of the type described above in connection with FIGS. 3 and 4.

FIGS. 6A and 6B depict a computer motherboard **600** adapted to use a variable-width memory in accordance with an embodiment of the invention.

FIG. 7 depicts a portion **700** of motherboard **600** detailing the signal-line configuration.

FIG. 8 depicts portion **700** of FIG. 7 with a memory module **800** and shorting module **810** installed.

FIGS. 9A–9D depict memory configurations that respectively accommodate one, two, three, or four memory modules.

FIG. 10 (prior art) shows a floor plan of a conventional 1 Gb DRAM **1000** having a 16-bit wide data path D0–D15.

FIGS. 11A and 11B depict a high-level floor plan of a DRAM **1100** featuring a configurable core.

FIG. 12 depicts a specific implementation of a configurable core **1200** and associated circuitry.

FIGS. 13A–F are simplified block diagrams of core **1200** of FIG. 12 illustrating access timing in a number of memory-access configurations.

#### DETAILED DESCRIPTION

FIG. 3 depicts a variable-width memory **300** in accordance with an embodiment of the invention. Memory **300** is similar to SDRAM **200** of FIG. 2, like-numbered elements being the same. Memory **300** differs from SDRAM **200**, however, in that the memory core organization changes with device width, resulting in reduced power usage for relatively narrow memory configurations. Also advantageous, reorganizing the core for relatively narrow memory widths increases the number of logical memory banks, and consequently reduces the likelihood of bank conflicts. Fewer conflicts means improved speed performance. These and other benefits of the invention are detailed below.

Much of the operation of memory **300** is similar to SDRAM **200** of FIG. 2. A discussion of those portions of memory **300** in common with SDRAM **200** is omitted here for brevity. The elements of FIG. 3 described above in connection with FIG. 2 are numbered in the two-hundreds (e.g., 2XX) for convenience. In general, the first digit of numerical designations indicates the figure in which the identified element is introduced.

Memory **300** includes a configurable memory core **305**. In the example, memory core **305** includes eight physical memory banks PB0–PB7, though the number of physical banks may vary according to need. Physical banks PB0–PB7 are interconnected such that they can be combined to form different numbers of logical banks. In the example, pairs of physical banks (e.g., PB0 and PB1) can be combined to form four logical banks LB0–LB3, collections of four physical banks (e.g., PB0–PB3) can be combined to form two logical banks LB4 and LB5, and all eight physical banks can be combined to form a single logical bank LB0–7. Assuming, for simplicity, that each physical bank PB0–PB7 includes a single data I/O terminal, memory core **305** can be configured as a one-bit-wide memory with eight logical banks, a two-bit-wide memory with four logical banks, a four-bit-wide memory with two logical banks, or an eight-bit-wide memory with one logical bank.

Some configuration logic **310** controls the configuration of memory core **305** via a data control circuit **315**. Configuration logic **310** also controls the data width through a collection of latches **320** and a collection of I/O buffers **325**. As detailed below, data control circuit **315** includes some



## 5

data routing logic, such as a crossbar switch, to provide flexible routing between the memory banks and data terminals DQs. The purpose and operation of these blocks is described below in more detail. As noted in FIG. 3, the data terminals (DQs) can be configured to have widths of x1, x2, x4, and x8.

FIG. 4A shows a specific implementation of a configurable core 400 and associated circuitry. In one embodiment, core 400 is a portion of memory 300 of FIG. 3. The number of physical banks is reduced to four physical banks PB0–PB3 in FIG. 4 for brevity. Memory 300 might include two memory “slices,” each of which comprises a memory core 400. The manner of extending the memory core of FIG. 4A to eight or more banks will be readily apparent to those of skill in the art.

The components of core 400 are similar to like-numbered elements in FIG. 3. For this embodiment, the serialization ratio is 1:1. Serialization ratios greater than 1:1 are possible with the addition of serial-to-parallel (write) and parallel-to-serial (read) conversion circuits. In this example, there are four physical banks PB0–3 supporting four read data bits and four write data bits. Generally, data control circuit 315 contains multiplexing logic for read operations and demultiplexing logic for write operations. The multiplexing logic and demultiplexing logic are designed to allow one, two, or four device data lines D0–D3 to be routed to the four physical banks PB0–PB3.

In the one-bit wide configuration, device data line D0 can be routed to/from any of the four physical banks PB0–PB3. In the 2-bit wide configuration (“x2”), device data lines D0 and D1 can be routed to/from physical banks PB0 and PB1 (collectively, logical bank LB0,1) or physical banks PB2 and PB3 (collectively logical banks LB2,3). Finally, in the 4-bit wide configuration, device data lines D0, D1, D2, and D3 can be routed to/from respective physical banks PB0, PB1, PB2, and PB3 (collectively, logical bank LB0–3). Core 400 can thus be configured as a one-, two-, or four-bank memory with respective widths of four (x4), two (x2), and one (x1) data bits.

Core 400 is a synchronous memory; consequently, each physical bank PB0–PB3 includes an input latch 405 and an output latch 410. Physical banks PB0–PB3 additionally include respective memory arrays MA0–MA3, sense amplifiers SA0–SA3, and bank-select terminals BS1–BS3. Asserting a bank select signal on one of terminals BS1–BS3 loads the data in the addressed location within the selected memory array into the respective one of sense amplifiers SA1–A3.

Latch 320 includes a pair of latches 415 and 420 for each physical bank PB0–PB3. Data control circuit 315 includes five multiplexers 425, 430, 435, 440, and 445 that communicate data between latch 320 and physical banks PB0–PB3. Multiplexers 425 and 430 are controlled by a write control signal WB; multiplexer 435 is controlled by a read control signal RA; multiplexer 440 is controlled by a write control signal WA; and multiplexer 445 is controlled by two read control signals RA and RB. Write control signals WA and WB and read control signals RA and RB are based on the selected data path width and bits of the requested memory address or transfer phase. Configuration logic 310 (FIG. 3) produces these signals in response to the programmed data width, whether the operation is a read or write operation, and appropriate addressing information.

Table 1 shows the control values used for data path slice widths of one, two, and four. Table 1 also indicates which of data terminals D0–D3 are used for each data width.

## 6

TABLE 1

WIDTH	WRITE		READ		DATA TERMINALS
	WA	WB	RA	RB	
1	1	1	A0	A1	D0
2	0	1	0	A0	D0 & D1
4	0	0	0	0	D0, D1, D2, & D3

When a width of one is selected during a read operation, the configuration logic 310 allows data from any one of the four physical banks PH0–PH3 to be presented at data terminal D0. Control signals RA and RB determine which data-bit signals will be presented at any given time. Control signals RA and RB are set (at this data width) to equal the two least-significant bits (A1, A0) of the memory address corresponding to the current read operation.

When a width of one is selected during a write operation, the circuit accepts the data bit signal from data terminal DQ0 and routes it to all four physical banks PB0–PB3 simultaneously. Control signals WA and WB are both set to a logical value of one to produce this routing. Other logic circuits (not shown) within configuration logic 310 control which of input latches 405 and 410 are active during any single write operation, so that each data bit signal is latched into the appropriate physical bank. For a given physical bank, only one of latches 405 and 410 is active during any given memory cycle.

When a width of two is selected during a read operation, configuration logic 310 allows two of the four data bit signals associated with physical banks PB0–PB3 to be present at data terminals DQ0 and DQ1. To obtain this result, control signal RA is set to 0, and control signal RB is equal to the lower bit (A0) of the memory address corresponding to the current read operation. Control signal RB determines which of two pairs of data bit signals (0 and 1 or 2 and 3) are presented at data terminals DQ0 and DQ1 during a given read operation.

When a width of two is selected during a write operation, configuration logic 310 accepts the data bit signals from physical banks PB0 and PB1 and routes them either to data terminals DQ0 and DQ1 or DQ2 and DQ3. In this configuration, physical banks PB0 and PB1 collectively form one logical bank LB0,1 and physical banks PB2 and PB3 collectively form a second logical bank LB2,3. Control signals WA and WB are set to 0 and 1, respectively, to obtain this result.

A width of four is selected by setting all of the control signals (RA, RB, WA, and WB) to 0. Read and write data signals are then passed directly between physical banks PB0–PB3 and corresponding data terminals DQ0–DQ3.

For each row access, data moves from memory arrays MA0–MA3 to their respective sense amplifiers SA0–SA3. Core 400 minimizes the power required to perform a row access by limiting each row access to the selected physical bank(s). To this end, bank-select signals on lines BS0–BS3 are only asserted to selected banks.

Configuration logic 310 determines which of physical banks PB0–PB3 are selected, and consequently which bank-select signals are asserted, based upon the selected device width and memory address. The following Table 2 summarizes the logic within configuration logic 310 that generates the appropriate bank-select signals.



TABLE 2

WIDTH	ADDRESS LINES A1:A0			
	00	01	10	11
1	BS0	BS1	BS2	BS3
2	BS0 & BS1	BS2 & BS3	BS0 & BS1	BS2 & BS3
4	BS0-BS3	BS0-BS3	BS0-BS3	BS0-BS3

When core **400** is configured to have a width of one, the two least-significant address bits **A0** and **A1** are decoded to select one of physical banks **PB0-PB3**; when core **400** is configured to have a width of two, address bit **A0** enables the physical banks within either of logical banks **LB0,1** or **LB2,3**; and when core **400** is configured to have a width of four, address bits **A0** and **A1** are ignored and all physical banks **PB0-PB3** are selected (i.e., enabled).

The circuit of FIG. **4A** is just one example of many possible designs. Other embodiments will benefit from other configurations. For example, it is possible to use more or less elaborate data routing schemes to account for the different connection needs for memory systems with more or fewer modules. Moreover, multiple memory cores **400** may be used to construct devices with greater than four device data connections. For example, a device having sixteen device data connections could use four memory cores while supporting three programmable widths; namely, 16, 8, or 4-bits widths. There are many possible alternatives for the number and width of physical and logical banks, the number of device data connections per device, serialization ratios, and data-path widths.

All data to and from memory core **400** passes through data line-DO terminal **DQ0** in the x1 mode, terminals **DQ0** and **DQ1** in the x2 mode, and terminals **DQ0-DQ3** in the x4 mode. FIG. **48** depicts an embodiment **450** that benefits from a more flexible routing scheme in which the data terminals **DQ0-DQ3** can be routed to different input/output pins of the memory module upon which core **305** is mounted. Embodiment **450** substitutes data control circuit **315** of FIG. **4A** with a more flexible crossbar switch **460**. In the depicted embodiment, the data terminals to and from physical bank **PB0** can be routed to any of data connections **DQ0-DQ3** in the x1 mode; the data terminals to and from physical banks **PB0** and **PB1** can be routed to either data connections **DQ0** and **DQ1** or data connections **DQ2** and **DQ3**, respectively, in the x2 mode; and the data terminals to and from physical banks **PB0-PB3** can be routed to data connections **DQ0-DQ3**, respectively, in the x4 mode. U.S. Pat. Nos. 5,530,814 and 5,717,871 describe various types of crossbar switches, and are incorporated herein by reference.

FIG. **5A** depicts a memory module **500** that includes four variable-width memories **502** of the type described above in connection with FIGS. **3**, **4A**, and **4B**. Module **500**, typically a printed circuit board, also includes a number of conductive traces **505** that convey data between the data pins (**3**, **2**, **1**, **0**) of memories **502** and corresponding module pins **510**. In FIG. **5A**, each memory **502** is configured to be one-bit wide, and the resulting four data bits are connected to four consecutive ones of pins **510**. The selected traces are identified as bold lines; the selected module pins are cross-hatched.

FIG. **5B** depicts the same memory module **500** of FIG. **5A**; unlike in FIG. **5A**, however, each memory **502** is configured to be two-bits wide, and the resulting eight data bits are connected to eight consecutive ones of pins **510**. The memory module **500** of FIG. **5B** is thus configured to be twice as wide (and half as deep) as the same module **500** of

FIG. **5A**. As in FIG. **5A**, the selected traces are identified as bold lines; the selected pins are crosshatched.

FIG. **5C** depicts the same memory module **500** of FIGS. **5A** and **5B**; unlike in FIGS. **5A** and **5B**, however, each memory **502** is configured to be four-bits wide, and the resulting sixteen data bits are connected to sixteen consecutive ones of pins **510**. The memory module **500** of FIG. **5C** is thus configured to be twice as wide (and half as deep) as the same module **500** of FIG. **5B** and four times as wide (and one fourth as deep) as the same memory module **500** of FIG. **5A**. Once again, the selected traces are identified as bold lines; the selected pins are crosshatched.

FIGS. **6A** and **6B** depict a computer motherboard (or system backplane) **600** adapted to use a variable-width memory in accordance with an embodiment of the invention. Motherboard **600** includes a memory controller **605** and a plurality of electrical receptacles or connectors **610** and **615**. The connectors are memory module sockets, and are configured to receive installable/removable memory modules **620** and **625**.

Each of memory modules **620** and **625** comprises a module backplane **630** and a plurality of integrated memory circuits **635**. Each memory module also includes first and second opposed rows of electrical contacts (module pins) **640** along opposite surfaces of its backplane. Only one row of contacts **640** is visible in FIG. **6A**. There are corresponding rows of connector contacts (not visible in FIG. **6A**) in each of connectors **610** and **615**.

A plurality of signal lines, or "traces," extend between memory controller **605** and electrical connectors **610** and **615** for electrical communication with memory modules **620** and **625**. More specifically, there are a plurality of sets of signal lines, each set extending to a corresponding, different one of connectors **610** and **615**. A first set of signal lines **645** extends to first electrical connector **610**, and a second set of signal lines **650** extends to second electrical connector **615**. Motherboard **600** also has a third set of signal lines **655** that extends between the two connectors.

In the embodiment shown, the signal lines comprise system data lines—they carry data that has been read from or that is to be written to memory modules **620** and **625**. It is also possible that other signal lines, such as address and control lines, would couple to the memory modules through the connectors. These additional signal lines could have a different interconnection topology than what is shown for signal lines **645**, **650**, and **655**.

The routing of the signal lines is more clearly visible in FIG. **6B**, in which memory modules **620** and **625** have been omitted for clarity. The illustrated physical routing is shown only as a conceptual aid—actual routing is likely to be more direct, through multiple layers of a printed-circuit board.

FIG. **7A** depicts a portion **700** of motherboard **600** detailing the signal-line configuration. This view shows cross-sections of connectors **610** and **615**. Electrical conductors, traces, and/or contacts are indicated symbolically in FIG. **7A** by relatively thick solid or dashed lines. Each of the three previously described sets of signal lines is represented by a single one of its conductors, which has been labeled with the reference numeral of the signal line set to which it belongs. The respective lines of a particular set of signal lines are routed individually in the manner shown.

As discussed above, each connector **610** and **615** has first and second opposed rows of contacts. FIG. **7A** shows individual contacts **705** and **710** corresponding respectively to the two contact rows of each connector. It is to be understood that these, again, are representative of the remaining contacts of the respective contact rows.



As is apparent in FIG. 7A, the first set of signal lines **645** extends to first contact row **705** of first connector **610**. The second set of signal lines **650** extends to the first contact row **705** of second connector **615**. In addition, a third set of signal lines **655** extends between the second contact row **710** of first connector **610** and second contact row **710** of second connector **615**. The third set of signal lines **655** is represented by a dashed line, indicating that these lines are used only in certain configurations; specifically, signal lines **655** are used only when a shorting module is inserted into connector **610** or **615**. Such a shorting module, the use of which will be explained in more detail below, results in both sets of signal lines **645** and **650** being configured for communications with a single memory module.

The system of FIG. 7A can be configured to include either one or two memory modules. FIG. 8 illustrates the first configuration, which includes a memory module **800** in the first connector **610** and a shorting module **810** in the second connector **615**. The shorting module has shorting conductors **815**, corresponding to opposing pairs of connector contacts, between the first and second rows of the second connector. Inserting shorting module **810** into connector **615** connects or couples the second set **650** of signal lines to the second contact row **705** of first connector **610** through the third set of signal lines **655**. In this configuration, the two sets of signal lines **645** and **650** are used collectively to communicate between memory controller **605** and memory module **800**.

In a two-module configuration, shorting bar **810** is replaced with a second memory module **800**. If modules **800** are adapted in accordance with the invention to support two width configurations and to include one half of the module pins **640** on either side, then there is no need for a switch matrix like data control circuit **315** of FIG. 4A or crossbar switch **460** of FIG. 4B. Instead, merely including shorting module **810** provides the memory controller access to the module pins **640** on both sides of the one module **800**. Alternatively, including two memory modules **800** will provide the memory controller access to the same half of the module pins **640** (those on the left-hand side of connector **610**) on both memory modules; the other half of the module pins **640** are not used. More complex routing schemes can likewise be employed to support additional modules and width configurations. The two-module configuration thus provides the same data width as the single-module configuration, with each module providing half the width.

For a more detailed discussion of motherboard **600**, see copending U.S. patent application Ser. No. 09/797,099 filed Feb. 28, 2001 entitled "Upgradeable Memory System with Reconfigurable Interconnect," by Richard E. Perego et al. which is incorporated herein by reference.

In some embodiments, the access configurations of the memory modules are controllable and programmable by memory controller **605** in the manner described above in connection with FIGS. 3, 4A, 4B, 5A, and 5B. In such embodiments, the memory controller may be adapted to detect which connectors have installed memory modules, and to set the configuration of each module accordingly. This allows either one or two memory module to be used in a system without requiring manual configuration steps. If one module is used, it may be configured to use two signal-line sets for the best possible performance. If two memory modules are present, they may each be configured to use one signal-line set. This idea can be extended to support memory systems that can accommodate more than two memory modules, though the routing scheme becomes more complex with support for additional modules.

The integrated memory circuit can be configured for the appropriate access mode using control pins. These control pins might be part of the signal line sets **645**, **650**, and **655**, or they might be part of a different set of signal lines. These control pins might be dedicated to this configuration function, or they might be shared with other functions. Also, the integrated memory circuit might utilize programmable fuses to specify the configuration mode. Integrated memory circuit configurability might also be implemented, for example, by the use of jumpers on the memory modules. Note that the memory capacity of a module remains the same regardless of how it is configured. However, when it is accessed through one signal line set it requires a greater memory addressing range than when it is accessed through two signal line sets. Also note that the two configurations shown in FIGS. 6–8 could also be implemented with a shorting connector instead of a shorting module. A shorting connector shorts its opposing contacts when no module is inserted (the same result as when the connector **615** of FIG. 7B has a shorting module inserted). A shorting connector with a memory module inserted is functionally identical to the connector **610** in FIG. 7.

As noted above, the general signal line scheme can be generalized for use with  $n$  connectors and memory modules. Generally stated, a system such as this uses a plurality of signal-line sets, each extending to a respective module connector. At least one of these sets is configurable or bypassable to extend to a connector other than its own respective connector. Stated alternatively, there are 1 through  $n$  sets of signal lines that extend respectively to corresponding connectors 1 through  $n$ . Sets 1 through  $n-1$  of the signal lines are configurable to extend respectively to additional ones of the connectors other than their corresponding connectors. FIGS. 9A–9D illustrate this generalization, in a memory system **900** in which  $n=4$ .

Referring first to FIG. 9A, this configuration includes a memory controller **905**; four memory slots or connectors **910**, **915**, **920**, and **925**; and four signal line sets **930**, **935**, **940**, and **945**. Each signal line set is shown as a single line, and is shown as a dashed line when it extends beneath one of the connectors without connection. Physical connections of the signal line sets to the connectors are shown as solid dots. Inserted memory modules are shown as diagonally hatched rectangles, with solid dots indicating signal connections. Note that each inserted memory module can connect to up to four signal line sets. The number of signal line sets to which it actually connects depends upon the connector into which it is inserted. The connectors are identical components, but appear different to the memory modules because of the routing pattern of the four signal line sets on the motherboard.

Each signal line set extends to a corresponding connector. Furthermore, signal line sets **935**, **940**, and **945** are extendable to connectors other than their corresponding connectors: signal line set **935** is extendable to connector **925**; signal line set **940** is extendable to both connectors **920** and **925**; signal line set **945** is extendable to connector **925**. More specifically, a first signal line set **930** extends directly to a first memory connector **925** without connection to any of the other connectors. It connects to corresponding contacts of the first contact row of connector **925**. A second signal line set **935** extends directly to a second memory connector **920**, where it connects to corresponding contacts of the first contact row. The corresponding contacts of the second contact row are connected to corresponding contacts of the first contact row of first connector **925**, allowing the second signal line set to bypass second connector **920** when a shorting module is placed in connector **920**.



## 11

A third signal line set **940** extends directly to a third memory connector **915**, where it connects to corresponding contacts of the first contact row. The corresponding contacts of the second contact row are connected to corresponding contacts of the first contact row of connector **920**. The corresponding second contact row contacts of connector **920** are connected to the corresponding contacts of the first contact row of connector **925**.

A fourth signal line set **945** extends directly to a fourth memory connector **910**, where it connects to corresponding contacts of the first contact row of connector **910**. The corresponding contacts of the second contact row are connected to corresponding contacts of the first contact row of first connector **925**.

This configuration, with appropriate use of shorting or bypass modules, accommodates one, two, three, or four physically identical memory modules. Each memory module permits simultaneous access through one, two, or four of its four available signal line sets. In the configuration of FIG. **9A**, a single memory module is inserted in first connector **925**. This memory module is configured to permit simultaneous accesses on all of its four signal line sets, which correspond to all four signal line sets. Connectors **910**, **915**, and **920** are shorted by inserted shorting modules as shown so that signal line sets **935**, **940**, and **945** extend to connector **925**.

FIG. **9B** illustrates a second configuration in which connectors **910** and **915** are shorted by inserting shorting modules. Thus, signal line sets **930** and **945** extend to connector **925** and the inserted memory module is configured to permit simultaneous accesses on these two signal line sets. Signal line sets **935** and **940** extend to connector **920** and the inserted memory module is configured to permit simultaneous accesses on these two signal line sets.

FIG. **9C** illustrates a third configuration in which connector **910** is shorted by inserting a shorting module, and memory modules are positioned in connectors **915**, **920**, and **925**. Signal line sets **930** and **945** extend to connector **925** and the inserted memory module is configured to permit simultaneous accesses on these two signal line sets. Signal line set **935** extends to connector **920** and the inserted memory module is configured to permit accesses on this signal line set. Signal line set **940** extends to connector **915** and the inserted memory module is configured to permit accesses on this signal line set.

FIG. **9D** illustrates a fourth configuration, with a memory module in each of the four available memory connectors. Each module is connected to use a respective one of the four signal line sets, with no shorting modules in use.

An interesting aspect of a memory device with programmable data access width relates to the characteristic of the device that its bandwidth may generally be reduced as its data width is narrowed. As device bandwidth is reduced, opportunities increase for altering the device's memory array configuration to provide greater independence between array partitions.

FIG. **10** shows an example of a conventional 1 Gb density DRAM **1000** with a 16-bit wide data path **D0–D15**. FIG. **10** shows a high-level floor plan of the DRAM die, including left (“L”) and right (“R”) bank subdivisions, row decoders, column decoders, I/O sense amps (I/O), and data pin locations **D0–D7** and **D8–D15**. A pair of regions **1005** and **1012** within memory banks **B0-L** and **B0-R** (i.e., the left and right halves of bank **0**) indicates a sample page location for an 8KB page within bank zero. 4KB worth of sense amp circuitry for the left and right halves of DRAM **1000** are accessed in parallel via a pair of multiplexers **1010** and **1015**

## 12

to form an 8KB page. In this design, data from left and right halves of the die are accessed in parallel to meet the device peak bandwidth requirement. This also allows the data paths for the left and right halves of the die to be largely independent. (This aspect of some embodiments is discussed in more detail below in connection with FIG. **12**.)

FIGS. **11A** and **11B** depict a high-level floor plan of a DRAM **1100** featuring a configurable core in accordance with one embodiment. DRAM **1100** can operate as DRAM **1000** of FIG. **10**, but can also be configured to reduce peak device bandwidth by a factor of two. Such a bandwidth reduction allows the full amount of device bandwidth to be serviced by either the left half (FIG. **11A**) or right half (FIG. **11B**) of the device. In this embodiment, the eight-active device data connections **D0–D7**—shown in bold—are located on the left side of the die, requiring that a data path **1105** be provided from the right side memory array to the left side data connections **D0–D7**. With the memory array divided into left-and right halves, it becomes feasible to manage banks on each side independently. In this case, the 16-bit wide device that supported eight independent banks accessed via data terminals **D0–D15** (like DRAM **1000** of FIG. **10**) can be reconfigured as an 8-bit wide device supporting 16 independent banks, with data access provided via either data terminals **D0–D7** or **D8–D15**.

There is typically some incremental circuit overhead associated with increasing the bank count of the device, setting a practical limit to the number of independent banks that could potentially be supported. However, a performance improvement related to the increased number of banks may justify some increase in device cost.

In the embodiment of FIGS. **11A** and **11B**, device page size is reduced for the 8-bit wide configuration (4KB) relative to the 16-bit wide configuration (8KB). Reducing the page size is attractive from a power consumption perspective because fewer sense amps are activated during a RAS operation. In addition to activating fewer sense amps, it is also possible to subdivide word lines using a technique known as “sub-page activation.” In this scheme, word lines are divided into multiple sections, one or more of which are activated for a particular RAS operation. This technique typically adds some incremental die area overhead in exchange for reduced power consumption and potentially improved array access or cycle times.

The examples highlighted in FIGS. **11A** and **11B** are intended to illustrate the concept of how a configurable array organization can be used to reduce power consumption and increase the number of logical memory banks. Write transactions are not described for this embodiment, although the same principles of power reduction and memory bank count apply to writes as well. The basic principles of configurable array organization can be exploited regardless of the type or capacity of memory device.

FIG. **12** depicts a specific implementation of a configurable core **1200** and associated circuitry, the combination of which may be integrated to form a memory component. Core **1200** is similar to core **450** of FIG. **4B**, like-named elements being the same. Core **1200** provides the same functionality as core **450**, but the configuration and switching logic is modified to afford users the ability to partition the four physical banks **PB0–PB3** into two separately addressable memories, each of which can be either one or two bits wide. Some elements are omitted from the depiction of FIG. **12** for brevity. For example, core **1200** may also include registers **405** and **410**.

Physical bank **PB0** includes a row decoder **RD0**, a memory array **MA0**, a sense amp **SA0** (actually a collection



of sense amplifiers), and a column decoder CD0. Each of the remaining physical banks PB1–PB3 includes identical structures. The row decoders, memory banks, sense amps, and column decoders are omitted from FIG. 4B for brevity, but are included in FIG. 12 to illustrate an addressing scheme that enables core 1200 to independently access address logical blocks LB0,1 and LB2,3.

Address buffers 225 and 230, introduced in FIG. 3, connect directly to the row and column decoders of physical banks PB2 and PB3. Configuration logic 310, also introduced in FIG. 3, connects to the bank-select terminals BS3–0 and to a crossbar switch 1207. Address buffers 225 and 230 are also selectively connected to the row and column decoders in physical banks PB0 and PB1 via a multiplexer 1205.

The configuration and switching logic of core 1200 is extended to include a second set of address buffers (row and column) 1209 and a second set of configuration logic 1210. Address buffers 1209 connect to the row and column decoders in physical banks PB0 and PB1 via multiplexer 1205. Configuration logic 1210 connects to crossbar switch 1207—the data control circuit in this embodiment—and to bank-select terminals BS0 and BS1 via multiplexer 1205. A configuration-select bus CONF from configuration logic 310 includes three control lines C0–C2 that connect to crossbar switch 1207. Line C2 additionally connects to the select terminal of multiplexer 1205. In this embodiment, mode register 220 (FIG. 3) is adapted to store configuration data establishing the levels provided on lines C0–C2.

Core 1200 supports four operational modes, or “configurations,” in addition to those described above in connection with FIGS. 3, 4A, and 4B. These modes are summarized below in Table 3.

TABLE 3

CONF#	CORE CONFIGURATION	C2	C1	C0
1	Single Address, Variable Width	1	X	X
2	Separate Addresses, separate 2-bit busses (DQ3/DQ2 and DQ1/DQ0)	0	0	0
3	Separate Addresses, memories share lines DQ1 and DQ0	0	0	1
4	Separate Addresses, memories share lines DQ3 and DQ2	0	1	0
5	Separate Addresses, banks configured to be 1-bit wide, data on lines DQ0 and DQ1	0	1	1

Core 1200 is operationally identical to core 450 of FIG. 4B if each of lines C0–C2 is set to logic one. In that case, the logic one on line C2 causes multiplexer 1205 to pass the address from address buffers 225 and 230 to physical banks PB0 and PB1. The logic levels on lines C0 and C1 are irrelevant in this configuration.

Driving line C2 to a voltage level representative of a logic zero causes multiplexer 1205 to convey the contents of the second set of address buffers 1209 to physical banks PB0 and PB1, and additionally causes crossbar switch 1207 to respond to the control signals on lines C0 and C1. Logical banks LB0,1 and LB2,3 are thereby separated to provide independent memory access. Logical banks LB0,1 and LB2,3 are separately addressable in each of configurations two through five of Table 3. Though not shown, logical banks LB0,1 and LB2,3 can be adapted to receive either the same clock signal or separate clock signals.

In configuration number two, crossbar switch 1207 accesses logical bank LB0,1 on lines DQ0 and DQ1 and

logical bank LB2,3 on lines DQ2 and DQ3. Core 1200 is therefore divided into a pair of two-bit memories accessed via separate two-bit data busses.

In configuration number three, crossbar switch 1207 alternatively accesses either logical bank LB0,1 or logical bank LB2,3 via lines DQ0 and DQ1. Core 1200 is therefore divided into two separately addressable two-bit memories that share a two-bit data bus. Configuration number four is similar, but access is provided via lines DQ2 and DQ3.

Configuration number five divides core 1200 into two separately addressable, one-bit-wide memories. In effect, each pair of physical blocks within logical blocks LB0,1 and LB2,3 is combined to form a single-bit memory with twice the address locations of a parallel configuration. Each of the resulting one-bit-wide memories is then separately accessible via one bus line.

The modes of Table 3 are not exhaustive. More control signals and/or additional control logic can be included to increase the available memory configurations. For example, configuration number five might be extended to include the ability to select the bus line upon which data is made available, or the two-bit modes could be extended to provide data on additional pairs of bus lines.

The mode-select aspect allows core 1200 to efficiently support data of different word lengths. Processors, which receive instructions and data from memory like core 1200, are sometimes asked to alternatively perform complex sets of instructions on relatively small data structures or perform simple instructions on relatively large data structures. In graphics programs, for example, the computationally simple task of refreshing an image employs large data structures, while more complex image processing tasks (e.g., texture mapping and removing hidden features) often employ relatively small data structures. Core 1200 can dynamically switch between configurations to best support the task at hand by altering the contents of mode register 220 (FIG. 3). In the graphics-program example, instructions that contend with relatively large data structures might simultaneously access both logical blocks LB0,1 and LB2,3 in parallel, and instructions that contend with relatively small data structures might access logical blocks LB0,1 and LB2,3 separately using separate addresses. Core 1200 may therefore provide more efficient memory usage. As with cores 400 and 450, core 1200 minimizes the power required to perform a row access by limiting each row access to the selected physical bank(s).

FIG. 13A is a simplified block diagram 1300 of core 1200 of FIG. 12 illustrating memory access timing in one memory-access mode. In this example, core 1200 is configured to deliver full-width data from combined logical blocks LB2,3 and LB0,1. The pairs of memory blocks within each logical block LB2,3 and LB0,1 are combined for simplicity of illustration. At time T1, the data stored in row address location ADD0 in each of logical blocks LB2,3 and LB0,1 are each loaded simultaneously into respective sense amplifiers SA2/3 and SA0/1. The row address ADD0 used for each logical block is the same. Then, at time T2, the contents at the same column address of the two sense amplifiers are accessed simultaneously with data lines DQ3/2 and DQ1/0 via switch 1207. Time T1 precedes time T2.

FIG. 13B is a block diagram 1310 of core 1200 of FIG. 12 illustrating access timing in a second memory-access mode. In this example, core 1200 is configured to alternatively deliver half-width data by separately accessing logical blocks LB2,3 and LB0,1. At time T1, the contents of row address ADD0 in logical block LB2,3 loads into sense amplifiers SA2/3. At another time T2 (where T2 may be



## 15

earlier or later than T1), the contents of row address ADD0 in local block LB0/1 loads into sense amplifiers SA0/1. Of interest, at each of times T1 and T2 only the accessed physical blocks are enabled using the appropriate bank-select signals BS3-0 (see FIG. 12). The content at a column address of sense amplifiers SA2/3 is accessed at time T3 via the data lines DQ0/1. The content at the same column address of sense amplifiers SA0/1 is accessed at another time T4 via the data lines DQ0/1 (where T4 may be earlier or later than T3). Time T1 precedes time T3, and time T2 precedes T4.

FIG. 13C is a simplified block diagram 1315 of core 1200 of FIG. 12 illustrating access timing in a third memory-access configuration. As in the example of FIG. 13A, core 1200 is configured to deliver full-width data from combined logical blocks LB2,3 and LB0,1; unlike the example of FIG. 13A, however, diagram 1315 illustrates the case in which logical blocks LB2,3 and LB0,1 are addressed separately. At time T1, the contents of row address ADD0 in logical block LB2,3 and row address ADD1 in logical block LB0,1 are loaded substantially simultaneously into respective sense amplifiers SA2/3 and SA0/1. The term “substantially simultaneous” is used here to indicate the possibility that these two operations are not precisely simultaneous (coincident), but nevertheless overlap. The content at a first column address of sense amplifiers SA2/3 is accessed at time T2 via the data lines DQ0/1. The content at a second column address of sense amplifiers SA0/1 is accessed substantially simultaneously at time T2 via the data lines DQ0/1. Time T1 precedes time T2.

FIG. 13D is a block diagram 1320 of core 1200 of FIG. 12 illustrating access timing in a fourth memory-access mode. With respect to timing, diagram 1320 is similar to diagram 1310 of FIG. 13B. Diagram 1320 differs from diagram 1310, however, in that each of logical block LB2,3 and LB0,1 is independently addressed. Core 1200 can therefore interleave data from different addresses in logical banks LB2,3 and LB0,1 and provide the resulting data on data lines DQ1 and DQ0. Specifically, at time T1, the contents of row address ADD0 in logical block LB2,3 loads into sense amplifiers SA2/3. At another time T2 (where T2 may be earlier or later or the same as T1), the contents of another row address ADD1 in logical block LB0/1 loads into sense amplifiers SA0/1 (ADD0 and ADD1 may be the same or different). The content at a first column address of sense amplifiers SA2/3 is accessed at time T3 via the data lines DQ0/1. The content at a second column address of sense amplifiers SA0/1 is accessed at another time T4 via the data lines DQ0/1 (where T4 may be earlier or later than T3). Time T1 precedes time T3, and time T2 precedes T4.

FIG. 13E is a simplified block diagram 1325 of core 1200 illustrating access timing in a mode that delivers full-width data from combined logical blocks LB2,3 and LB0,1. With respect to timing, diagram 1325 is similar to diagram 1300 of FIG. 13A. Diagram 1325 differs from diagram 1300, however, in that each of logical blocks LB2,3 and LB0,1 is independently addressed.

FIG. 13F is a simplified block diagram 1330 of core 1200 illustrating access timing in a mode that delivers half-width data from independently addressed logical blocks LB2,3 and LB0,1. The flow of data in diagram 1330 is similar to that of diagram 1320 of FIG. 13D. However, diagram 1330 differs from diagram 1320 with respect to timing because the contents of address locations ADD0 of logical block LB2,3 and ADD1 of logical block LB0,1 are delivered to respective sense amplifiers SA2/3 and SA0/1 substantially simultaneously.

## 16

Although details of specific implementations and embodiments are described above, such details are intended to satisfy statutory disclosure obligations rather than to limit the scope of the following claims. Thus, the invention as defined by the claims is not limited to the specific features described above. Rather, the invention is claimed in any of its forms or modifications that fall within the proper scope of the appended claims, appropriately interpreted in accordance with the doctrine of equivalents.

What is claimed is:

1. A memory system comprising:

- a. a memory bus having a plurality of system data lines; and
- b. at least one memory module having:
  - i. a plurality of module pins connected to respective ones of the system data lines;
  - ii. a plurality of memory banks;
  - iii. configuration and control logic connected to the memory banks and adapted to configure the memory banks to form one of several memory-width configurations; and
  - iv. a switch matrix connected between the memory banks and the plurality of module pins, the switch matrix adapted to programmably connect the memory banks to all or a subset of the system data pins;

c. wherein doubling the number of memory modules in the system automatically halves the memory-width configuration for each module.

2. A memory system comprising:

- a. a first address bus;
- b. a second address bus;
- c. a first memory bank, including a first address port connected to the first address bus, a first memory array, and a first sense amplifier;
- d. a second memory bank, including a second address port connected to the first and second address busses, a second memory array, and a second sense amplifier; and
- e. configuration and control logic connected to the memory banks and adapted to configure the memory system in one of at least two addressing modes, including a first addressing mode and a second addressing mode;
- f. wherein the memory system moves first data from a first address in the first memory array to the first sense amplifier and moves second data from a second address in the second memory array to the second sense amplifier in the first addressing mode, the first and second addresses being selected using the first address bus; and
- g. wherein the memory system moves third data from a third address in the first memory array to the first sense amplifier and moves fourth data from a fourth address in the second memory array to the second sense amplifier in the second addressing mode, wherein the third and fourth addresses are separate addresses selected using the respective first and second address busses.

3. The memory system of claim 2, further comprising a memory access bus including first and second data lines, wherein the first and second data are provided substantially simultaneously on the respective first and second data lines and the third and fourth data are provided substantially on the respective first and second data lines.

4. The memory system of claim 2, further comprising a memory access bus including first and second data lines,



17

wherein the first and second data are provided substantially simultaneously on the respective first and second data lines and the third data is provided on the first data line before the fourth data is provided on the second data line.

**5.** A memory component comprising:

- a. a first memory bank having a first memory array and a first sense amplifier;
- b. a second memory bank having a second memory array and a second sense amplifier;
- c. configuration logic connected to the first and second memory banks, the configuration logic supporting at least first and second memory-access modes;
  - i. wherein the configuration logic loads first data from the first memory array into the first sense amplifier and substantially simultaneously loads second data from the second memory array into the second sense amplifier in the first memory-access mode; and
  - ii. wherein the configuration logic sequentially loads the first data from the first memory array into the first sense amplifier and the second data from the second memory array into the second sense amplifier in the second memory-access mode.

**6.** The memory component of claim **5**, further comprising first and second data terminals and data control circuitry, wherein the data control circuitry is adapted to transfer the first and second data substantially simultaneously in the respective first and second sense amplifiers to the first and second data terminals in the first memory-access mode.

**7.** The memory component of claim **5**, further comprising first and second data terminals and data control circuitry, wherein the data control circuitry is adapted to sequentially transfer the first and second data in the respective first and second sense amplifiers to the first data terminal in the second memory-access mode.

**8.** The memory component of claim **5**, wherein the first data is stored at a first address in the first memory array, the second data is stored at a second address in the second memory array, and wherein the first and second addresses differ.

**9.** A memory component comprising:

- a. first and second data terminals;
- b. first and second memory arrays;
- c. configuration and switching logic connected between the data terminals and the memory arrays, the configuration and switching logic supporting at least first and second memory-access modes;

18

- i. wherein the configuration and switching logic, substantially simultaneously, transfers first data from the first memory array to the first data terminal and transfers second data from the second memory array to the second data terminal in the first memory-access mode; and
- ii. wherein the configuration and switching logic sequentially transfers the first data from the first memory array to the first data terminal and the second data from the second memory array to the first data terminal in the second memory-access mode.

**10.** The memory component of claim **9**, further comprising a first memory output latch disposed between the first memory array and the configuration and switching logic and a second memory output latch disposed between the second memory array and the configuration and switching logic.

**11.** The memory component of claim **10**, further comprising a first sense amplifier disposed between the first memory output latch and the first memory array and a second sense amplifier disposed between the second memory output latch and the second memory array.

**12.** The memory component of claim **11**, wherein the configuration and switching logic transfers, substantially simultaneously, the first data from the first memory array to the first sense amplifier and transfers the second data from the second memory array to the second sense amplifier before transferring, substantially simultaneously, the first data from the first sense amplifier to the first data terminal and transferring the second data from the second sense amplifier to the second data terminal in the first memory-access mode.

**13.** The memory component of claim **11**, wherein the configuration and switching logic transfers, substantially simultaneously, the first data from the first memory array to the first sense amplifier and transfers the second data from the second memory array to the second sense amplifier before sequentially transferring the first data from the first sense amplifier to the first data terminal and transferring the second data from the second sense amplifier to the second data terminal in the first memory-access mode.

**14.** The memory component of claim **10**, further comprising a mode register connected to the configuration and switching logic and adapted to define whether the memory component operations in the first memory-access mode of the second memory-access mode.

\* \* \* \* \*